

(19) World Intellectual Property Organization
International Bureau(43) International Publication Date
10 April 2003 (10.04.2003)

PCT

(10) International Publication Number
WO 03/030461 A1

(51) International Patent Classification: H04L 12/46

(21) International Application Number: PCT/US02/29074

(22) International Filing Date:
12 September 2002 (12.09.2002)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/966,349 28 September 2001 (28.09.2001) US

(71) Applicant: INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US).

(72) Inventors: HOOPER, Donald; 19 Main Circle, Shrewsbury, MA 01545 (US). HIRNAK, Stephanie; 11 Newbury Lane, Bedford, NH 03110 (US).

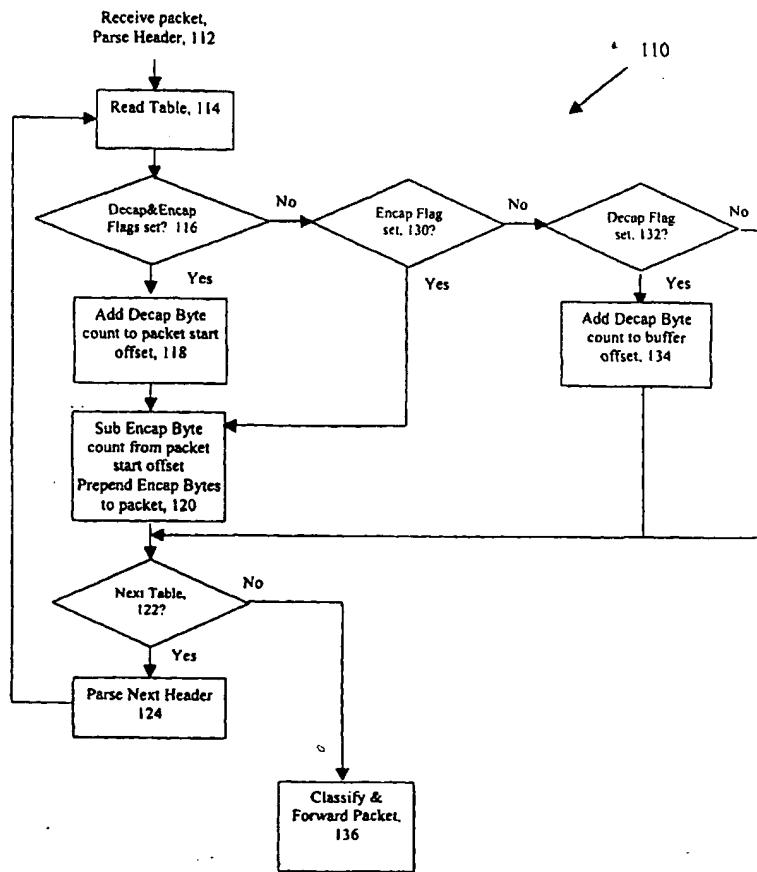
(74) Agents: HARRIS, Scott, C.; Fish & Richardson, P.C., 4350 La Jolla Village Drive, Suite 500, San Diego, CA 92122 et al. (US).

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: METHOD, APPARATUS AND COMPUTER PROGRAM FOR THE DECAPSULATION AND ENCAPSULATION OF PACKETS WITH MULTIPLE HEADERS



(57) Abstract: A method of forwarding a network packet is described. The method includes reading a table (114) containing a plurality of flags to determine which of the plurality of flags is set or cleared (116, 130, 132) and performing an operation on the packet to decapsulate (118, 120) or encapsulate (134) the packet in accordance with values of the flags.

WO 03/030461 A1



Published:

— *with international search report*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

METHOD, APPARATUS AND COMPUTER PROGRAM FOR THE DECAPSULATION AND ENCAPSULATION OF PACKETS WITH MULTIPLE HEADERS

BACKGROUND

This invention relates to forwarding network packets
5 between network domains.

Packets are routed through a series of router devices, each of which stores and forwards packets on its way from a source to a destination. For example, a packet may start out as an Internet packet, be forwarded over an ATM
10 (asynchronous transfer mode path) and then back to Ethernet onto a corporate network to its final intended recipient. As the network passes through these network domains, various header encapsulations may be added to or removed from the packet. Some connections use point-to-point protocol (PPP)
15 whereas others use multiprotocol label switching MPLS, layer to tunneling protocol LTTP, ATM and so forth.

DESCRIPTION OF DRAWINGS

FIG. 1 is a block diagram of a communication system employing a hardware based multithreaded processor.

20 FIGS. 2-1 to 2-4 are a detailed block diagram of a hardware based multithreaded processor of FIG. 1.

FIG. 3 is a block diagram depicting a functional arrangement of the multithreaded processor of FIG. 2.

25 FIG. 4 is a block diagram depicting data structures in memory used for the processor of FIG. 1.

FIG. 5 is a block diagram depicting formats for forwarding tables used in the tables of FIG. 4.

FIG. 6 is a flow chart depicting a generic packet forwarding process.

30 FIG. 7 is a flow chart depicting an alternative aspect of the packet forwarding process of FIG. 6.

DETAILED DESCRIPTION

Referring to FIG. 1, a communication system 10 includes a parallel, hardware-based multithreaded processor 12. The hardware-based multithreaded processor 12 is coupled to a 5 bus such as a PCI bus 14, a memory system 16 and a second bus 18. The system 10 is especially useful for tasks that can be broken into parallel subtasks or functions.

Specifically hardware-based multithreaded processor 12 is useful for tasks that are bandwidth oriented rather than 10 latency oriented. The hardware-based multithreaded processor 12 has multiple microengines 22 each can be simultaneously active and work on multiple hardware controlled threads that independently work on a task.

The hardware-based multithreaded processor 12 also 15 includes a central controller 20 that assists in loading microcode control for other resources of the hardware-based multithreaded processor 12 and performs other general purpose computer type functions such as handling protocols, exceptions, extra support for packet processing where the 20 microengines pass the packets off for more detailed processing such as in boundary conditions. In one embodiment, the processor 20 is a Strong Arm® based architecture. The general purpose microprocessor 20 has an operating system. Through the operating system the 25 processor 20 can call functions to operate on microengines 22a-22f. The processor 20 can use any supported operating system preferably a real time operating system such as, MicrosoftNT real-time, VXWorks.

The hardware-based multithreaded processor 12 also 30 includes a plurality of function microengines 22a-22f. Functional microengines (microengines) 22a-22f each maintain a plurality of program counters in hardware and states associated with the program counters. Effectively, a

corresponding plurality of sets of threads can be simultaneously active on each of the microengines 22a-22f while only one is actually operating at any one time.

In one embodiment, there are six microengines 22a-22f as shown. Each microengines 22a-22f has capabilities for processing four hardware threads. The six microengines 22a-22f operate with shared resources including memory system 16 and bus interfaces 24 and 28. The memory system 16 includes a Synchronous Dynamic Random Access Memory (SDRAM) controller 26a and a Static Random Access Memory (SRAM) controller 26b. SDRAM memory 16a and SDRAM controller 26a are typically used for processing large volumes of data, e.g., processing of network payloads from network packets. The SRAM controller 26b and SRAM memory 16b are used in a networking implementation for low latency, fast access tasks, e.g., accessing look-up tables, memory for the core processor 20, and so forth.

The six microengines 22a-22f access either the SDRAM 16a or SRAM 16b based on characteristics of the data. Thus, low latency, low bandwidth data is stored in and fetched from SRAM, whereas higher bandwidth data for which latency is not as important, is stored in and fetched from SDRAM. The microengines 22a-22f can execute memory reference instructions to either the SDRAM controller 26a or SRAM controller 16b.

Advantages of hardware multithreading can be explained by SRAM or SDRAM memory accesses. As an example, an SRAM access requested by a Thread_0, from a microengine will cause the SRAM controller 26b to initiate an access to the SRAM memory 16b. The SRAM controller controls arbitration for the SRAM bus, accesses the SRAM 16b, fetches the data from the SRAM 16b, and returns data to a requesting microengine 22a-22b. During an SRAM access, if the

microengine e.g., 22a had only a single thread that could operate, that microengine would be dormant until data was returned from the SRAM. By employing hardware context swapping within each of the microengines 22a-22f, the

5 hardware context swapping enables other contexts with unique program counters to execute in that same microengine. Thus, another thread e.g., Thread_1 can function while the first thread, e.g., Thread_0, is awaiting the read data to return. During execution, Thread_1 may access the SDRAM memory 16a.

10 While Thread_1 operates on the SDRAM unit, and Thread_0 is operating on the SRAM unit, a new thread, e.g., Thread_2 can now operate in the microengine 22a. Thread_2 can operate for a certain amount of time until it needs to access memory or perform some other long latency operation, such as making

15 an access to a bus interface. Therefore, simultaneously, the processor 12 can have a bus operation, SRAM operation and SDRAM operation all being completed or operated upon by one microengine 22a and have one more thread available to process more work in the data path.

20 The hardware context swapping also synchronizes completion of tasks. For example, two threads could hit the same shared resource e.g., SRAM. Each one of these separate functional units, e.g., the FBUS interface 28, the SRAM controller 26a, and the SDRAM controller 26b, when they

25 complete a requested task from one of the microengine thread contexts reports back a flag signaling completion of an operation. When the flag is received by the microengine, the microengine can determine which thread to turn on.

One example of an application for the hardware-based

30 multithreaded processor 12 is as a network processor. As a network processor, the hardware-based multithreaded processor 12 interfaces to network devices such as a media access controller device e.g., a 10/100BaseT Octal MAC 13a

or a Gigabit Ethernet device 13b. In general the network process can interface to any type of communication device or interface that receives/sends large amounts of data. Communication system 10 functioning in a networking 5 application could receive a plurality of network packets from the devices 13a, 13b and process those packets in a parallel manner. With the hardware-based multithreaded processor 12, each network packet can be independently processed. Another example for use of processor 12 is a 10 print engine for a postscript processor or as a processor for a storage subsystem, i.e., RAID disk storage. A further use is as a matching engine. In the securities industry for example, the advent of electronic trading requires the use of electronic matching engines to match orders between 15 buyers and sellers. These and other parallel types of tasks can be accomplished on the system 10.

The processor 12 includes a bus interface 28 that couples the processor to the second bus 18. Bus interface 28 in one embodiment couples the processor 12 to the so- 20 called FBUS 18 (FIFO bus). The FBUS interface 28 is responsible for controlling and interfacing the processor 12 to the FBUS 18. The FBUS 18 is a 64-bit wide FIFO bus, which is currently gaining acceptance as the best bus for Media Access Controller (MAC) devices.

25 The processor 12 includes a second interface e.g., a PCI bus interface 24 that couples other system components that reside on the PCI 14 bus to the processor 12. The PCI bus interface 24, provides a high speed data path 24a to memory 16 e.g., the SDRAM memory 16a. Through that path 30 data can be moved quickly from the SDRAM 16a through the PCI bus 14, via direct memory access (DMA) transfers. Additionally, the PCI bus interface 24 supports target and master operations. Target operations are operations where

slave devices on bus 14 access SDRAMs through reads and writes that are serviced as a slave to target operation. In master operations, the processor core 20 sends data directly to or receives data directly from the PCI interface 24.

5 Each of the functional units are coupled to one or more internal buses. The processor includes an AMBA bus that couples the processor core 20 to the memory controller 26a, 26c and to an AMBA translator 30 described below. The processor also includes a private bus 34 that couples the 10 microengine units to SRAM controller 26b, AMBA translator 30 and FBUS interface 28. A memory bus 38 couples the memory controller 26a, 26b to the bus interfaces 24 and 28 and memory system 16 including flashrom 16c used for boot operations and so forth.

15 Referring to FIGS. 2-1 to 2-4, each of the microengines 22a-22f includes an arbiter that examines flags to determine the available threads to be operated upon. Any thread from any of the microengines 22a-22f can access the SDRAM controller 26a, SDRAM controller 26b or FBUS interface 28.

20 The memory controllers 26a and 26b each include a plurality of queues to store outstanding memory reference requests. The queues either maintain order of memory references or arrange memory references to optimize memory bandwidth. For example, if a thread_0 has no dependencies or relationship 25 to a thread_1, there is no reason that thread 1 and 0 cannot complete their memory references to the SRAM unit out of order. The microengines 22a-22f issue memory reference requests to the memory controllers 26a and 26b. The microengines 22a-22f flood the memory subsystems 26a and 26b 30 with enough memory reference operations such that the memory subsystems 26a and 26b become the bottleneck for processor 12 operation.

If the memory subsystem 16 is flooded with memory requests that are independent in nature, the processor 12 can perform memory reference sorting. Memory reference sorting improves achievable memory bandwidth. Memory reference sorting, as described below, reduces dead time or a bubble that occurs with accesses to SRAM. With memory references to SRAM, switching current direction on signal lines between reads and writes produces a bubble or a dead time waiting for current to settle on conductors coupling the SRAM 16b to the SRAM controller 26b.

That is, the drivers that drive current on the bus need to settle out prior to changing states. Thus, repetitive cycles of a read followed by a write can degrade peak bandwidth. Memory reference sorting allows the processor 12 to organize references to memory such that long strings of reads can be followed by long strings of writes. This can be used to minimize dead time in the pipeline to effectively achieve closer to maximum available bandwidth. Reference sorting helps maintain parallel hardware context threads. On the SDRAM, reference sorting allows hiding of pre-charges from one bank to another bank. Specifically, if the memory system 16b is organized into an odd bank and an even bank, while the processor is operating on the odd bank, the memory controller can start precharging the even bank. Precharging is possible if memory references alternate between odd and even banks. By ordering memory references to alternate accesses to opposite banks, the processor 12 improves SDRAM bandwidth.

The FBUS interface 28 supports Transmit and Receive flags for each port that a MAC device supports, along with an Interrupt flag indicating when service is warranted. The FBUS interface 28 also includes a controller 28a that performs header processing of incoming packets from the FBUS

18. The controller 28a extracts the packet headers and performs a microprogrammable source/destination/protocol hashed lookup (used for address smoothing) in SRAM. If the hash does not successfully resolve, the packet header is 5 sent to the processor core 20 for additional processing. The FBUS interface 28 supports the following internal data transactions:

10 FBUS unit (via AMBA bus) to/from processor Core.

FBUS unit (via private bus) to/from SRAM Unit.

FBUS unit (via Mbus) to/from SDRAM.

15 The FBUS 18 is a standard industry bus and includes a data bus, e.g., 64 bits wide and sideband control for address and read/write control. The FBUS interface 28 provides the ability to input large amounts of data using a series of input and output FIFO's 29a-29b. From the FIFOs 20 29a-29b, the microengines 22a-22f fetch data from or command the SDRAM controller 26a to move data from a receive FIFO in which data has come from a device on bus 18, into the FBUS interface 28. The data can be sent through memory controller 26a to SDRAM memory 16a, via a direct memory 25 access. Similarly, the microengines can move data from the SDRAM 26a to interface 28, out to FBUS 18, via the FBUS interface 28.

30 Data functions are distributed amongst the microengines. Connectivity to the SRAM 26a, SDRAM 26b and FBUS 28 is via command requests. A command request can be a memory request or a FBUS request. For example, a command request can move data from a register located in a microengine 22a to a shared resource, e.g., an SDRAM

location, SRAM location, flash memory or some MAC address. The commands are sent out to each of the functional units and the shared resources. However, the shared resources do not need to maintain local buffering of the data. Rather, 5 the shared resources access distributed data located inside of the microengines. This enables microengines 22a-22f, to have local access to data rather than arbitrating for access on a bus and risk contention for the bus. With this feature, there is a 0 cycle stall for waiting for data 10 internal to the microengines 22a-22f.

The data buses, e.g., AMBA bus 30, SRAM bus 34 and SDRAM bus 38 coupling these shared resources, e.g., memory controllers 26a and 26b are of sufficient bandwidth such that there are no internal bottlenecks. Thus, in order to 15 avoid bottlenecks, the processor 12 has an bandwidth requirement where each of the functional units is provided with at least twice the maximum bandwidth of the internal buses. As an example, the SDRAM can run a 64 bit wide bus at 83 MHz. The SRAM data bus could have separate read and 20 write buses, e.g., could be a read bus of 32 bits wide running at 166 MHz and a write bus of 32 bits wide at 166 MHz. That is, in essence, 64 bits running at 166 MHz which is effectively twice the bandwidth of the SDRAM.

The core processor 20 also can access the shared 25 resources. The core processor 20 has a direct communication to the SDRAM controller 26a to the bus interface 24 and to SRAM controller 26b via bus 32. However, to access the microengines 22a-22f and transfer registers located at any 30 of the microengines 22a-22f, the core processor 20 access the microengines 22a-22f via the AMBA Translator 30 over bus 34. The AMBA translator 30 can physically reside in the FBUS interface 28, but logically is distinct. The AMBA Translator 30 performs an address translation between FBUS

microengine transfer register locations and core processor addresses (i.e., AMBA bus) so that the core processor 20 can access registers belonging to the microengines 22a-22c.

The processor core 20 includes a RISC core 50
5 implemented in a five stage pipeline performing a single cycle shift of one operand or two operands in a single cycle, provides multiplication support and 32 bit barrel shift support. This RISC core 50 is a standard Strong Arm® architecture but it is implemented with a five stage
10 pipeline for performance reasons. The processor core 20 also includes a 16 kilobyte instruction cache 52, an 8 kilobyte data cache 54 and a prefetch stream buffer 56. The core processor 20 performs arithmetic operations in parallel with memory writes and instruction fetches. The core
15 processor 20 interfaces with other functional units via the ARM defined AMBA bus. The AMBA bus is a 32-bit bi-directional bus 32.

Referring to FIG. 3, the multiprocessor 12 is shown performing network routing functions. In one example, an
20 asynchronous transfer mode (ATM), Ethernet and other types of packets enter through the network interface MAC devices and are sent to the network processor 12. These packets are processed in an application on the general purpose microprocessor 20 or on another processor that is coupled
25 through the PCI bus interface (not shown). For reception and transmission of such packets, the application running on that processor 20 or the processor coupled through the PCI bus, makes use of a network stack 72, which includes network management, control and signaling processes 74 to manage
30 network communications.

The network stack 72 and the application run in the processor 20 that controls the microengines, or another processor coupled to the PCI bus. The paths of receive,

transmit and data forwarding represent the transport of the packets through the processor 12. The management control, signaling, and the network stack 72 usually are not involved in data forwarding. Essentially, the processor 20 receives 5 and transmits. The processor 20 generates new packets that are transmitted over the network. The processor 20 can be involved in data forwarding in the exceptional case. This would involve very unusual packets, which may need special handling and complex processing.

10 For data forwarding processes, the microengines 22a-22f are used. In some instances, data forwarding may occur at the general purpose processor 20 level. The signals Init is programmer's interface for initialization of microengine code. The signal Fini is used for termination (to put 15 control info in a known state). The microengines 22a-22f provide fast, store and forward capabilities. The engines use a multilayer generic look-up process that performs validation, classification, policing and filtering using parallel hardware supported threads of the process.

20 Exceptions and control packets are passed to the processor 20 for processing at the network stack 72. A ternary network stack (not shown) can be located off-chip at a host via the PCI port or device port. This can be used to off-load the processor 20 or centralized management and control 25 for one place. In some embodiments, the microengine is a compact RISC processor and can have limited instruction space. For this reason and for other reasons, it is desirable to reduce instruction code size when running multiple protocols. The network processor 12 implements a 30 generic forwarding process that can be used to handle various protocol types (both existing and future types) without exceeding instruction storage limits.

Referring now to FIG. 4, an management arrangement 80 for forwarding table structures 90 that are stored in memory is shown. The forwarding table structure management 80 includes a control and management structure 82 including a 5 network stack interface 84 and table managers 86. The table managers 86 manage routing tables 90 that are stored in SRAM and can include a plurality of tables such as shown in FIG. 4 including a layer 4 connection table 92, a layer 3 destination table 94, a layer 2 bridge table 96 and a layer 10 2 connection table 98. Additionally, data structures stored in memory can include a packet buffer 100, which is stored in DRAM. The microengines acting as packet data forwarding processors retrieve information from the routing tables 90 in SRAM and store and forward the packet information from 15 the packet buffer in DRAM. The multiple tables 90 are set up by the control management processor 20. For example, a layer 2 connection table 96 can be used for ATM virtual circuits, frame relay connections MPLS labels or other low level connections. A layer 2 bridge table 96 could be used 20 for Ethernet bridging. A layer 3 destination table 94 could be used for Internet protocol (IP) forwarding based on a destination IP address. The layer 4 connection table 92 could be used for IP forwarding based on source and destination ports, addresses and protocol. All these tables 25 may require that the packet be decapsulated or encapsulated.

Once the tables 90 are populated with forwarding information in a generally conventional manner, packet data forwarding processors can receive packets, perform table look-ups to obtain information and convert packets as 30 required by the table entry. The control management process sets up the tables 90 with a common format for the purpose of decapsulation and encapsulation.

Referring now to FIG. 5, exemplary table entries, a subset of which are included in each of the tables 90, is shown. The table entries include the following fields:

5

Forwarding Table Format

Decap Flag. Indicates whether the bytes should be stripped from the packet. If this flag is asserted, then the number of bytes to strip is in Decap Byte Count field.

10

Decap To Layer. This field specifies decapsulation of header layers up to the specified layer. The length of the layer and hence the decapsulation is determined by parsing the packet header.

15

Decap Byte Count. This field specifies the number of bytes to remove from the front of the packet. Decap is performed by adjusting the packet start offset in the packet buffer.

20

Current Encap. This field specifies an identifier of the current packet encapsulation type.

25

Encap Flag. Indicates whether bytes should be prepended to the packet. If this flag is asserted, then the number of bytes is in Encap Byte Count field, and the bytes to be encapsulated is in the Encap Header field.

Encap Byte Count. Number of bytes to be prepended to the packet.

30

Encap Header. The actual bytes to be prepended.

Next Table Type. If non-zero, this indicated that a further lookup is required. This gives the table type. For example, layer 3

routing or layer 4 connection table type. A layer 3 routing lookup would use a longest prefix match lookup algorithm using the destination IP address. A layer 4 connection 5 lookup would use a 104 bit hash algorithm using source and destination addresses, source and destination ports, and protocol.

10 **Next Table Addr.** There can be multiple next tables, and multiple next tables of the same type. This field specifies the base address of the table.

15 The flags get set or cleared by the management process. Signaling and setting up connections are part of the network system that will determine that a certain path through the network requires a change of the header. There can be many reasons why a header can change. Usually a header change is used when the protocol changes from one network domain to another.

20 Referring now to FIG. 6, a process 110 for encapsulating/decapsulating generic protocols is shown. Initially one of the microengines 22a-22f receives 112 a packet from the network interface. The packet is comprised of one or more headers followed by a payload. The 25 microengine, e.g., 22a copies the payload portion of the packet to a packet buffer in DRAM and it may place the packet at an offset in the buffer to make room for any new header that could be prepended to the packet for packet forwarding. The packet offset parameter for that packet is 30 set to a default value determined at the offset into the buffer. The microengine reads 114 in the first header of the packet and performs a layer 2 look-up. The layer 2 look-up will read the table layer 2 bridge table and/or

layer 2 connection table. The tables will return various parameters such as decap or encapsulation flags. The process 110 will determine 116 if the decap or encapsulation flags are set. If the decap and encapsulation flags are set, the process will add 118 5 the decap byte count to the packet start offset and will subtract 120 the encapsulation byte count from the packet start offset and prepend the encapsulation bytes to the packet. The process 110 tests 122 if there is a next table to examine by looking at the blank field in the currently read table. If 10 there is a next table, the process 110 will parse the next header 124, fetch and read the next table. The process 110 continues looking to test the decap or encapsulation flags being set.

If, however, the process did not determine that the 15 decap and encapsulation flags were set (116, above), it would determine 130 if the encapsulation flag or the decap flag were set 132. If the encapsulation flag was set, it will subtract 120 the encapsulation flag byte count from the start offset and prepend the encapsulation bytes to the packet. On the other hand, if the decap 20 flag was only set 132, the process will add 134 a decap byte count to the buffer offset and, in any event, will check the next table 112. When the process determines that it is at the end of checking the tables, it will then classify and forward 136 the packet in a conventional manner. That is, 25 the "no" condition indicates that the process can classify and forward. Forwarding the header can have the microengine take the header and send it to the processor 20 or elsewhere, so that it can get reassembled with the payload. Forwarding the header could also involve forwarding the 30 packet, etc.

Referring now to FIG. 7, in addition to specifying byte dissemination counts obtained from the look-up table. The look-up table may have the decap to layer field set in the

table. This field specifies that the front portion of a packet should be decapsulated up to a certain layer. As known, packets are defined in protocol layers used in the OSI (Open Systems Interconnect) seven layer network 5 protocol. After passing through physical layer 1, the first software layer seen by the network processor layer is layer 2, also referred to as the link layer. The length of the bytes to be decapsulated is determined by parsing the packet layers prior to the layer that is to be the new start of the 10 packet. The length can be added to the packet start offset.

FIG. 7 shows a variation where the decapsulation length is not specified in the table, but is determined by reading the packet itself. In other words, this would be a set of routines that would be inserted into the processing of FIG. 15 6 substitute encapsulated byte count from the packet into the offset.

A process 140 to determine this offset is shown in FIG. 7. The process 140 includes reading the table 142, determining that the decap to packet layer bit 144 has been 20 set, and if set, retrieve the length of the layer to be removed by parsing the header 146 and adding the length to the packet start offset 148. If the decap layer has not been set then the process simply skips. In any event, this process can be prepended to the process described in 25 conjunction with FIG. 6.

A typical use of a decap to layer bit is to specify a decapsulation up to the layer 3 IP header. If the packet encapsulation is a multiprotocol over an ATM network such as the RFC 1483 standard, the layer 2 header length is 30 determined by parsing the layer to header itself using the RFC 1483 length rules. However, if the packet encapsulation is classical IP the layer 2 length is determined by following the classical IP layer length rules. The packet

encapsulation may be known by the port type it came in on from the prepended custom header from that port or may be obtained from the first look-up table in the current encapsulation field.

5 Rather than each network protocol defining a separate protocol conversion this technique provides a generic approach. The approach saves code space and software development time-to-market. In an alternative embodiment, this technique can be implemented as a software library

10 routine, e.g., a generic software building block for decapsulation/encapsulation, where customers can insert their proprietary header encapsulation and a customer's vendor need not get involved with customer's proprietary protocol designs.

15 A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. Accordingly, other embodiments are within the scope of the following claims.

20

WHAT IS CLAIMED IS:

1. A method of forwarding a network packet comprises:
reading a table containing a plurality of flags to
determine which of the plurality of flags is set or cleared;
5 and
performing an operation on the packet to decapsulate or
encapsulate the packet in accordance with values of the
flags.
2. The method of claim 1 wherein the tables are populated
10 with forwarding information.
3. The method of claim 1 wherein the forwarding table
structures include a control and management structure
including a network stack interface and table managers.
4. The method of claim 1 wherein the table managers manage
15 routing tables and can include a plurality of tables
including a layer 4 connection table, a layer 3 destination
table, a layer 2 bridge table and a layer 2 connection
table.
5. The method of claim 1 wherein the tables include a flag
20 to indicates whether the bytes should be stripped from the

packet and a field that indicates the number of bytes to be stripped.

6. The method of claim 1 wherein the tables include a field that specifies decapsulation of header layers up to 5 the specified layer.

7. The method of claim 1 wherein the tables includes a field that specifies an identifier of the current packet encapsulation type.

8. The method of claim 1 wherein the tables include a flag 10 that indicates whether bytes should be prepended to the packet, a field that specifies the number of bytes and the bytes to be encapsulated.

9. The method of claim 1 wherein the tables include a Next 15 Table Type field which indicates that a further lookup is required and identifies the table type.

10. A method for encapsulating/decapsulating packets comprises:

receiving a packet;
reading in a first header of the packet and perform a 20 layer 2 look-up reading a connection table which return parameters;

determine if the table returned a decap or encapsulation flag.

11. The method of claim 10 wherein if the decap and encapsulation flags are set,

5 adding the decap byte count to a packet start offset and subtracting the encapsulation byte count from the packet start offset; and

prepending the encapsulation bytes to the packet.

12. The method of claim 10 further comprising:

10 determining if there is a next table to examine by looking at the blank field in the currently read table.

13. The method of claim 12 wherein if there is a next table,

parsing the next header and fetch and read the next table.

15 14. The method of claim 11 wherein if the decap and encapsulation flags were not set,

determine if the encapsulation flag or the decap flag were set.

15. The method of claim 11 wherein if the encapsulation flag was set,

20 subtract the encapsulation flag byte count from the start offset and prepend the encapsulation bytes to the packet.

16. The method of claim 11 wherein if the decap flag was set add a decap byte count to the buffer offset and check the next table.

17. The method of claim 11 wherein the packet is comprised
5 of one or more headers followed by a payload, the method further comprises:

copying the payload portion of the packet to a packet buffer.

18. The method of claim 17 wherein copying may place the
10 packet at an offset in the buffer to make room for any new header that could be prepended to the packet for packet forwarding.

19. A computer program product residing on a computer readable media for forwarding a network packet comprises
15 instructions to cause a computer to:

read a table containing a plurality of flags to determine, which of the plurality of flags is set or cleared; and

20 perform an operation on the packet to decapsulate or encapsulate the packet in accordance with values of the flags.

20. The computer program product of claim 19 wherein the tables are populated with forwarding information.
21. The computer program product of claim 19 wherein the forwarding table structures include a control and management structure including a network stack interface and table managers.
22. A computer program product residing on a computer readable media for forwarding a network packet comprises instructions to cause a computer to:
 - 10 receive a packet;
read in a first header of the packet and perform a layer 2 look-up reading a connection table which return parameters;
 - determine if the table returned a decap or encapsulation flag.
- 15 23. The computer program product of claim 22 wherein if the decap and encapsulation flags are set, the computer program executes instructions to:
add the decap byte count to a packet start offset and subtracting the encapsulation byte count from the packet start offset; and
20 prepend the encapsulation bytes to the packet.

24. The computer program product of claim 22 further comprising instructions to:

 determine if there is a next table to examine by looking at the blank field in the currently read table.

5 25. The computer program product of claim 24 wherein if there is a next table, the computer program executes instructions to:

 parse the next header and fetch and read the next table.

10 26. The computer program product of claim 22 wherein the packet is comprised of one or more headers followed by a payload, the computer program product further executes instructions to:

 copy the payload portion of the packet to a packet buffer.

27. The computer program product of claim 26 wherein instructions to copy place the packet at an offset in the buffer to make room for any new header that could be prepended to the packet for packet forwarding.

20 28. A processor for processing a network packet comprises:

a computer storage media storing instructions to cause a computer to:

read a table containing a plurality of flags to determine, which of the plurality of flags is set or 5 cleared; and

perform an operation on the packet to decapsulate or encapsulate the packet in accordance with values of the flags.

29. The processor of claim 28 wherein the table 10 contains forwarding information.

30. A method of decapsulating a network packet comprises:

reading a table containing a plurality of flags to determine which of the plurality of flags is set or cleared; 15 and

performing a decapsulate operation on the packet in accordance with values of the flags.

31. The method of claim 30 wherein the table contains forwarding information.

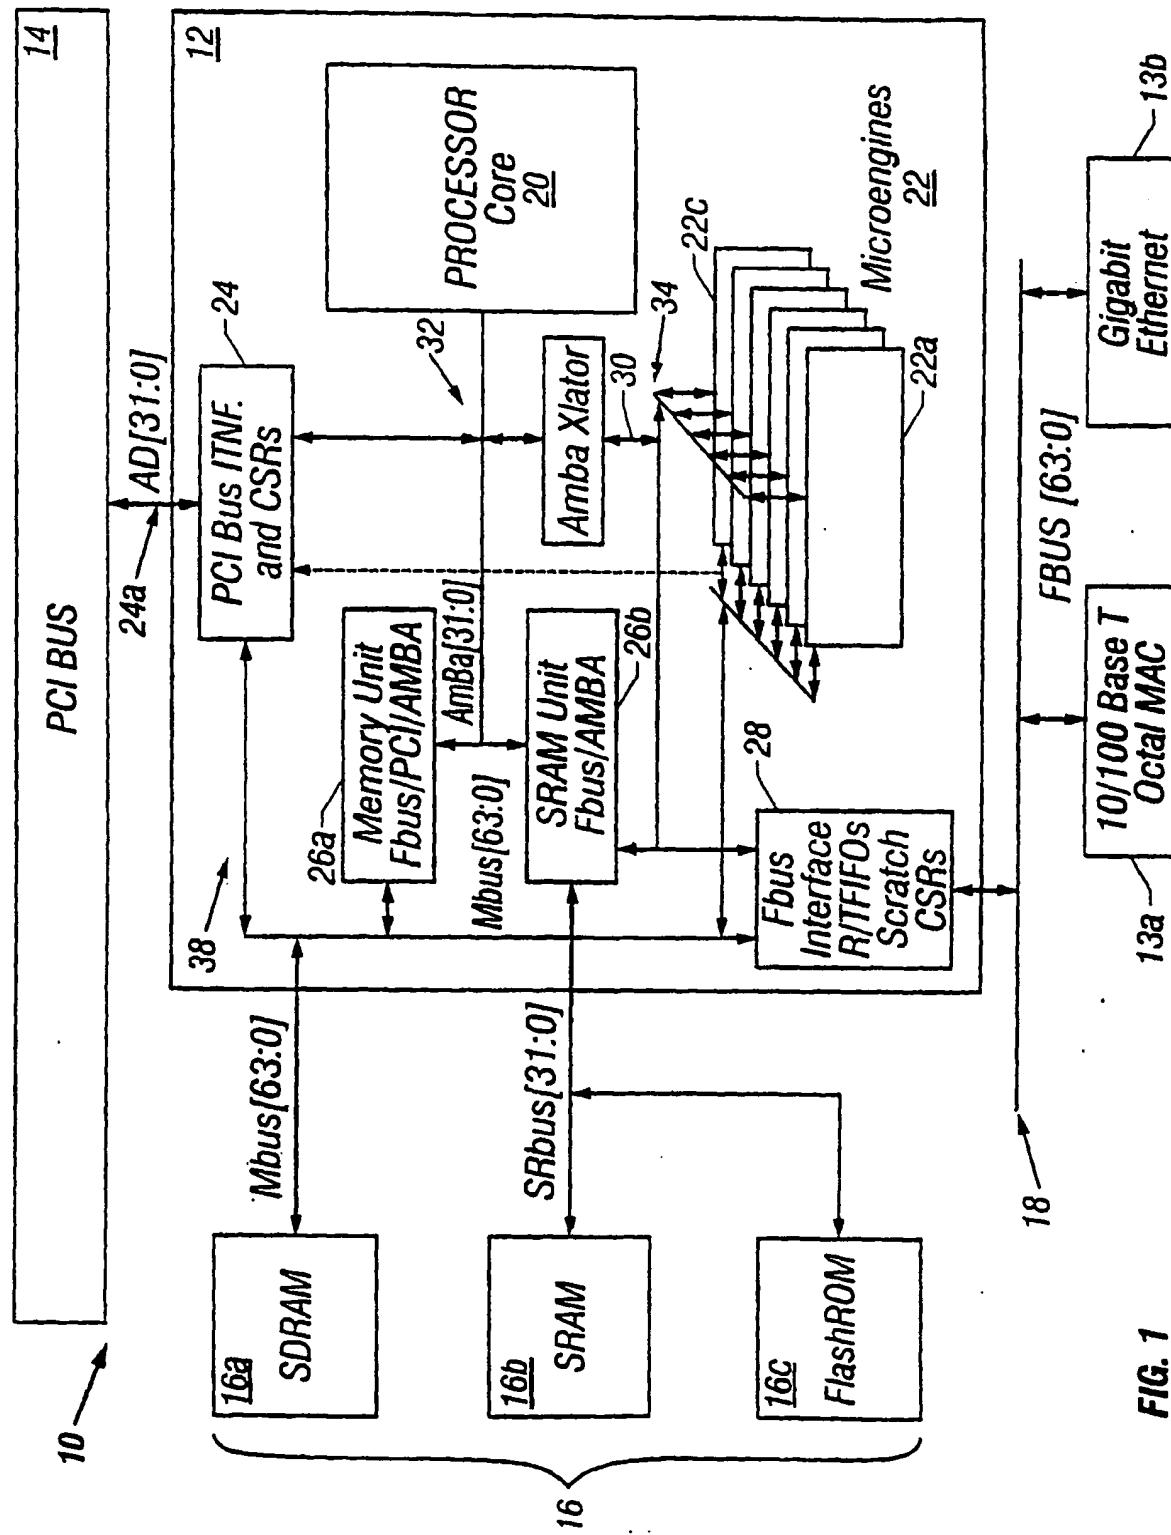


FIG. 1

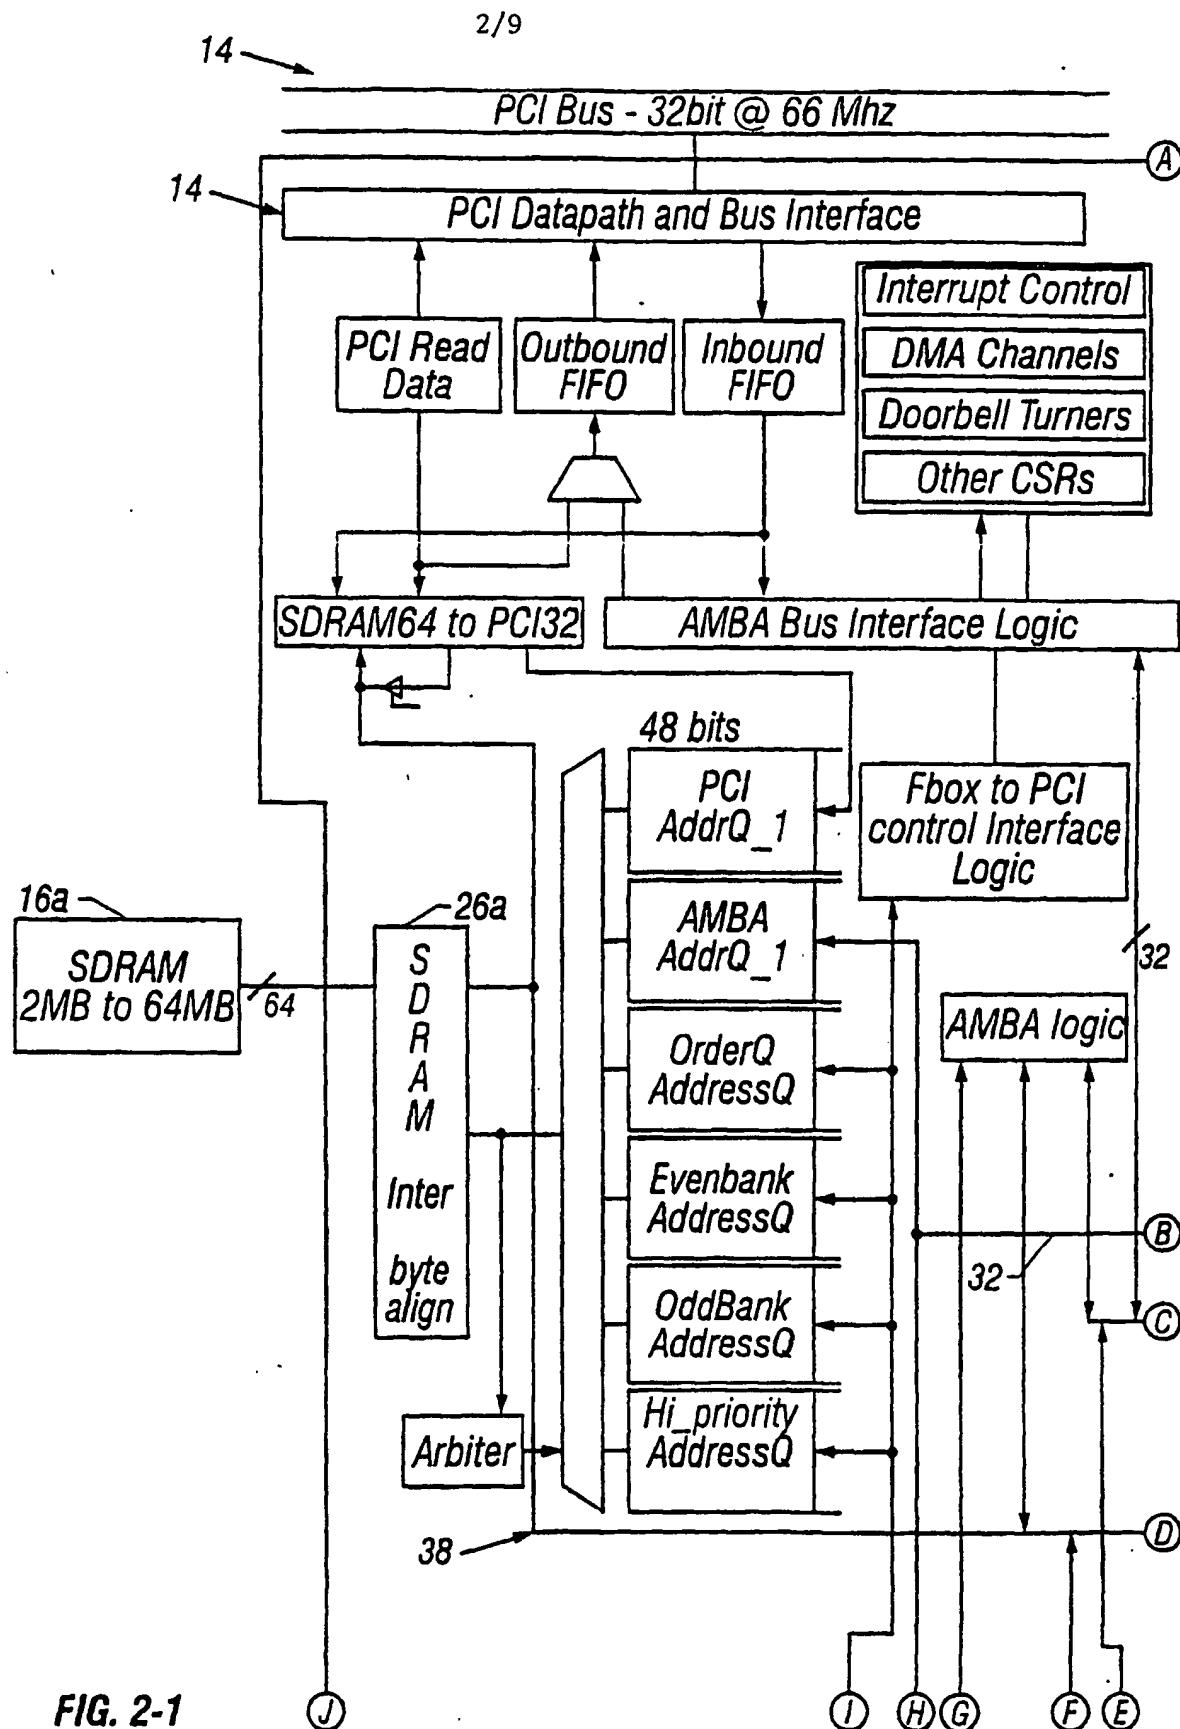
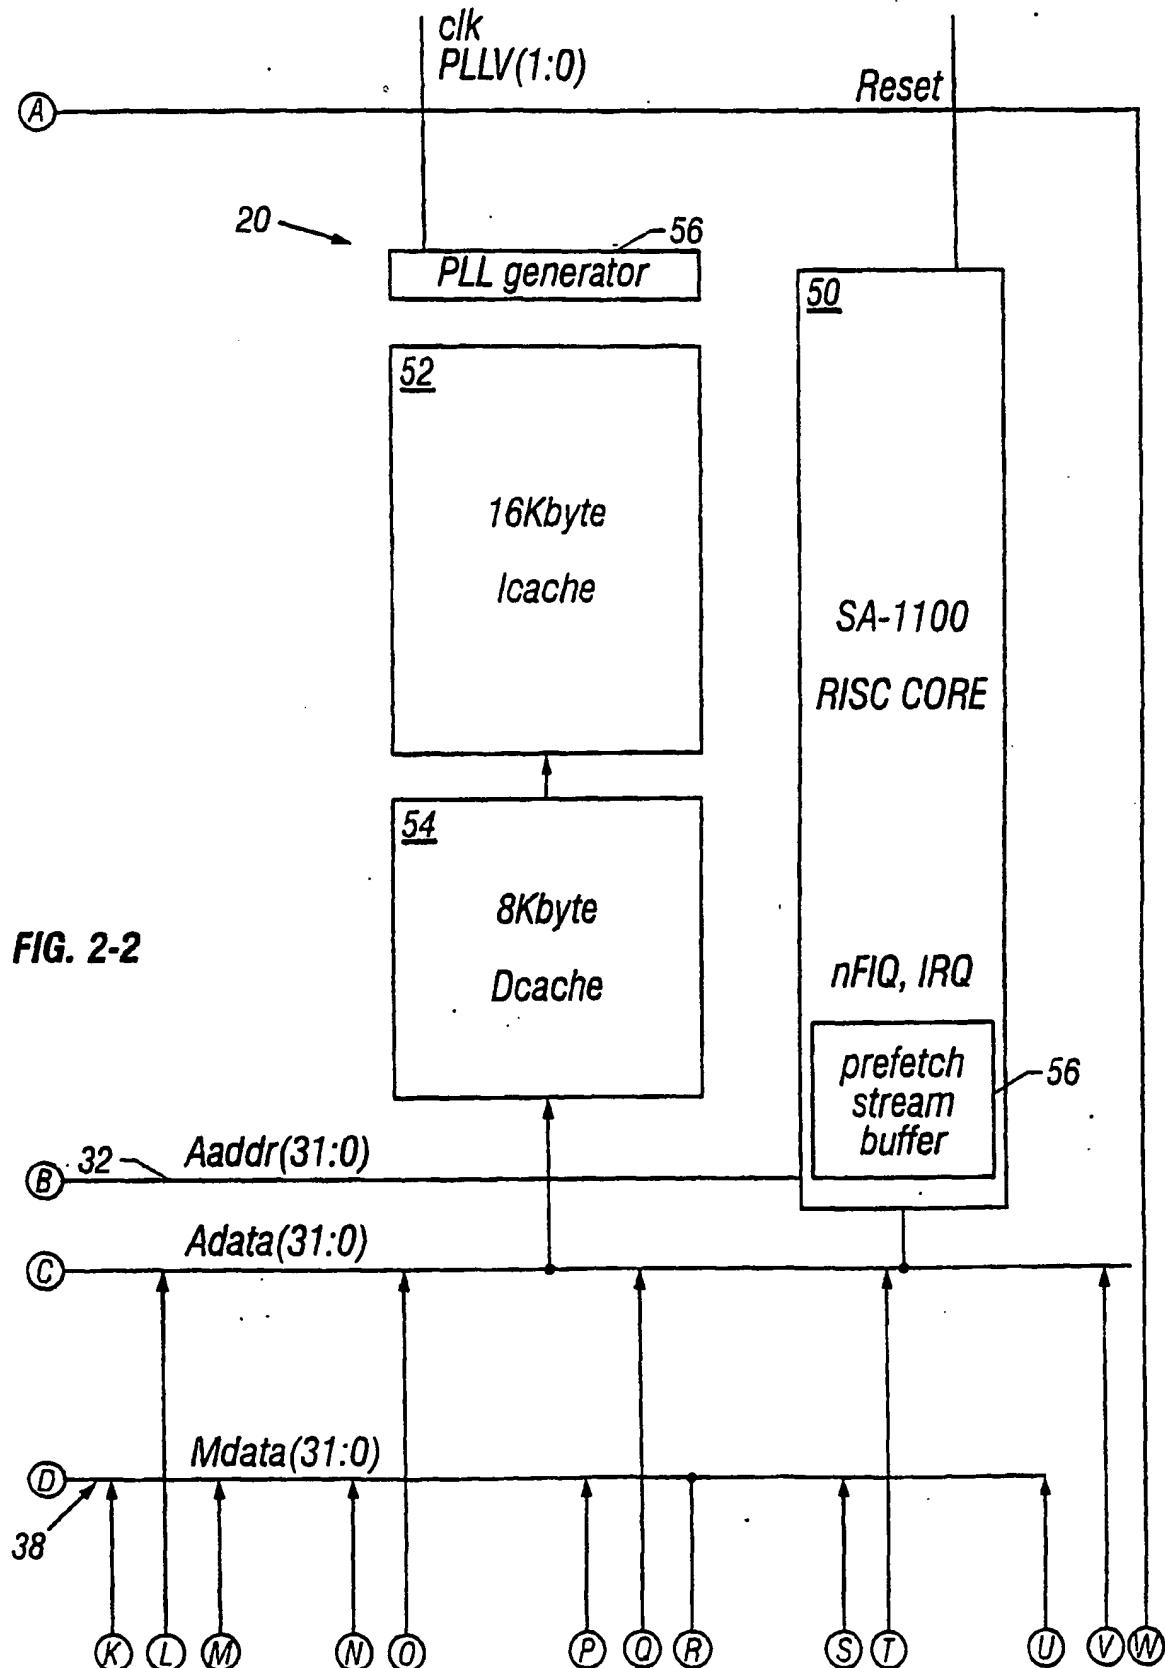


FIG. 2-1



4/9

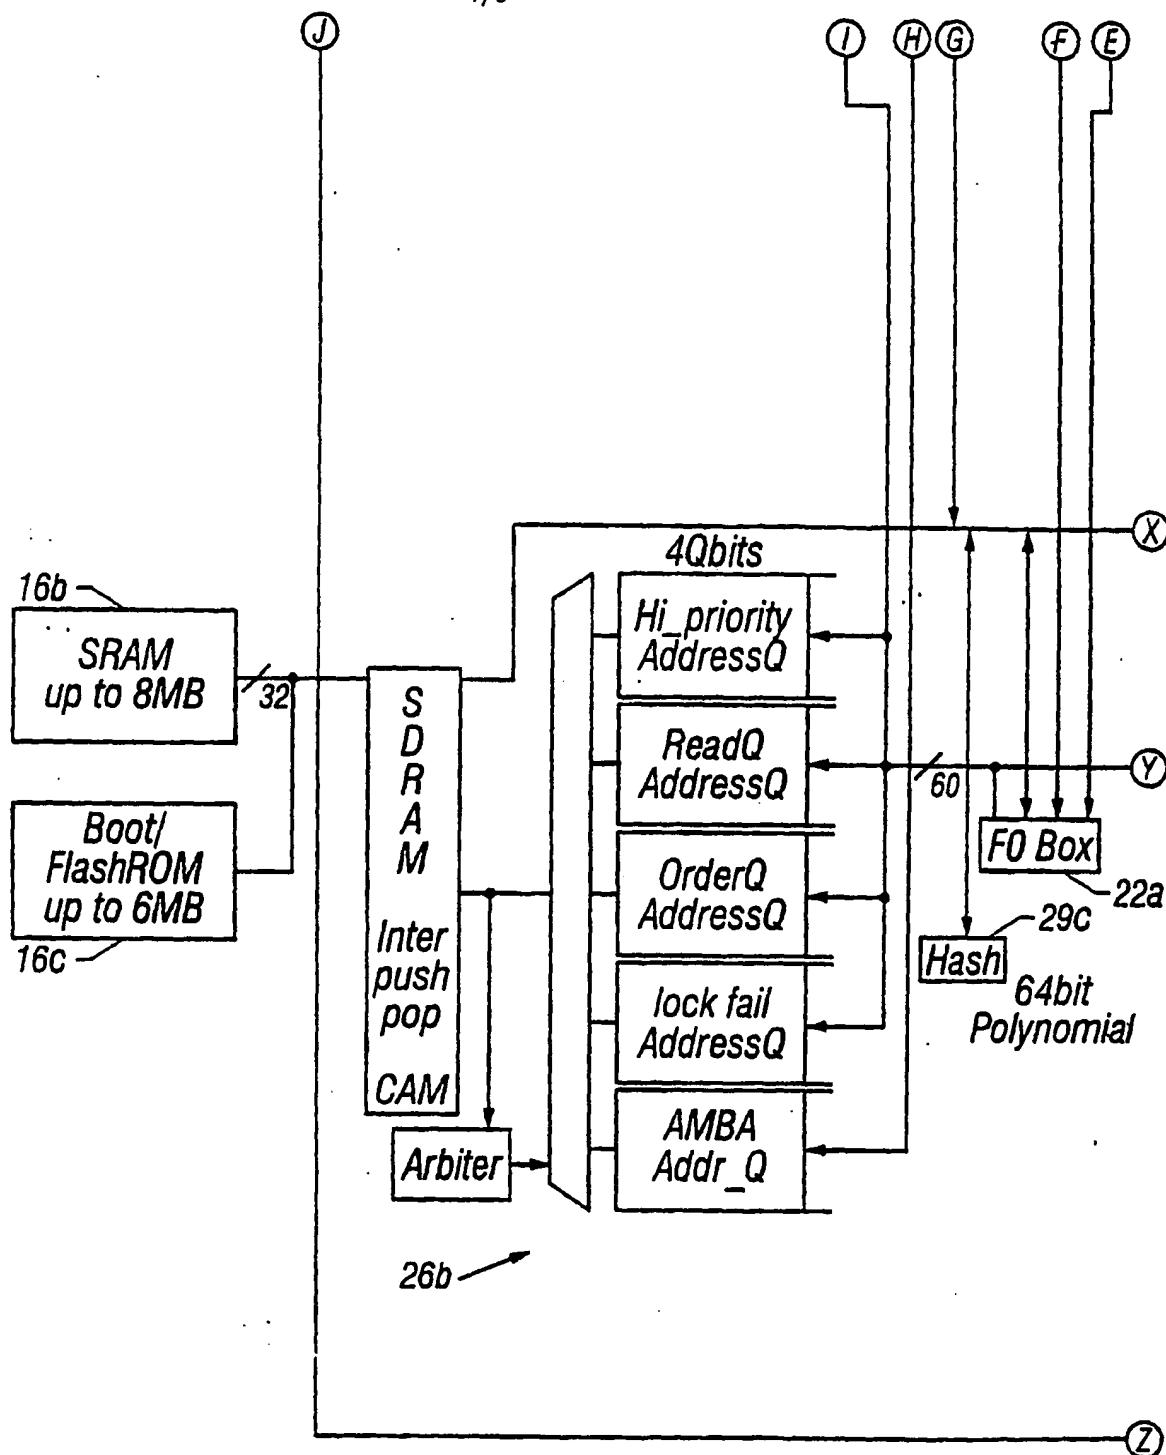


FIG. 2-3

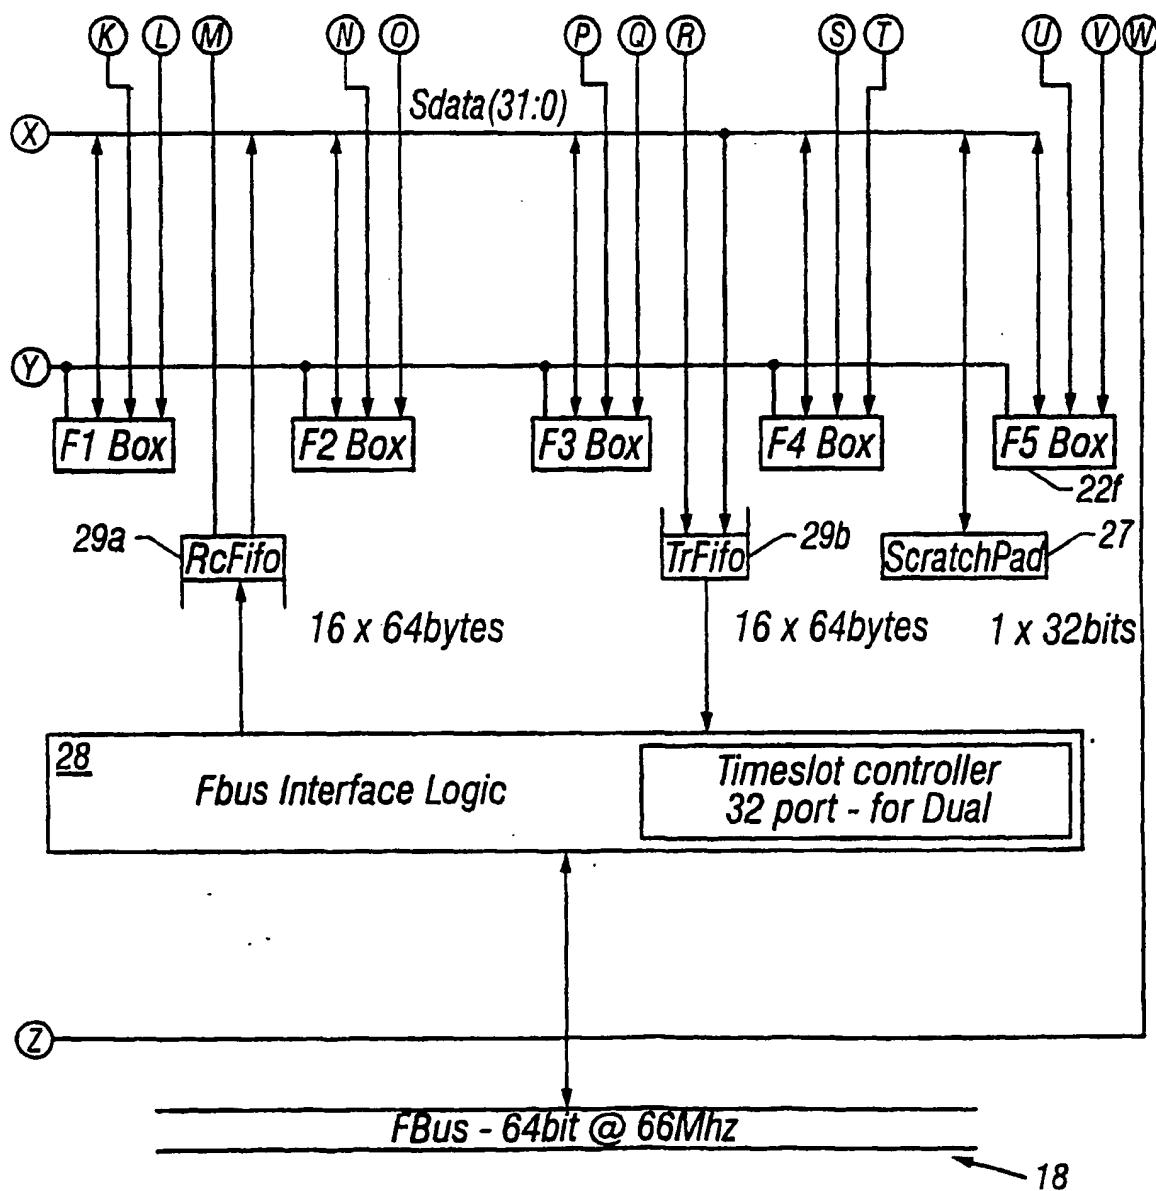
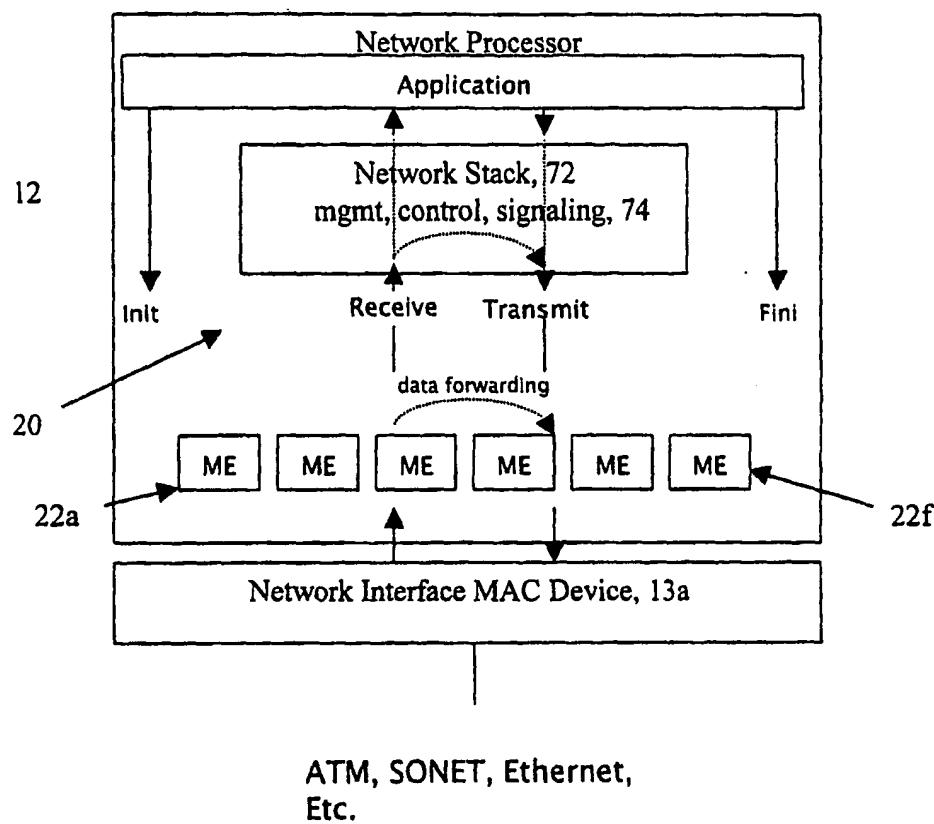


FIG. 2-4

FIG. 3



7/9

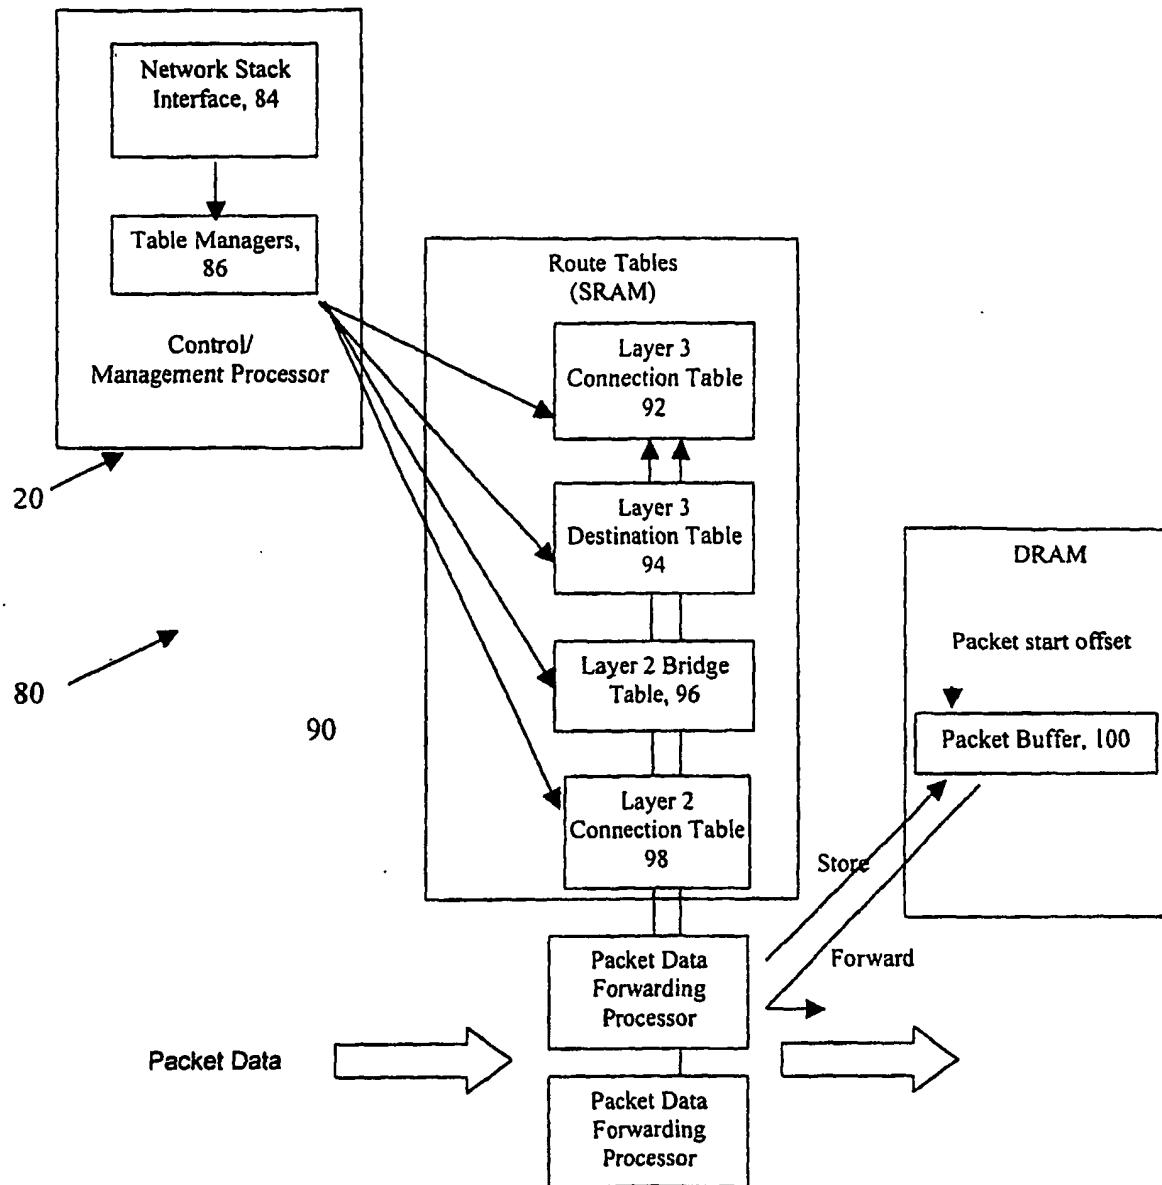


FIG. 4

Decap Flag	Decap To Layer	Decap Byte Count	Current Encap	Encap Flag	Encap Byte Count	Encap Bytes	Next Table Type	Next Table Addr
------------	----------------	------------------	---------------	------------	------------------	-------------	-----------------	-----------------

FIG. 5

8/9

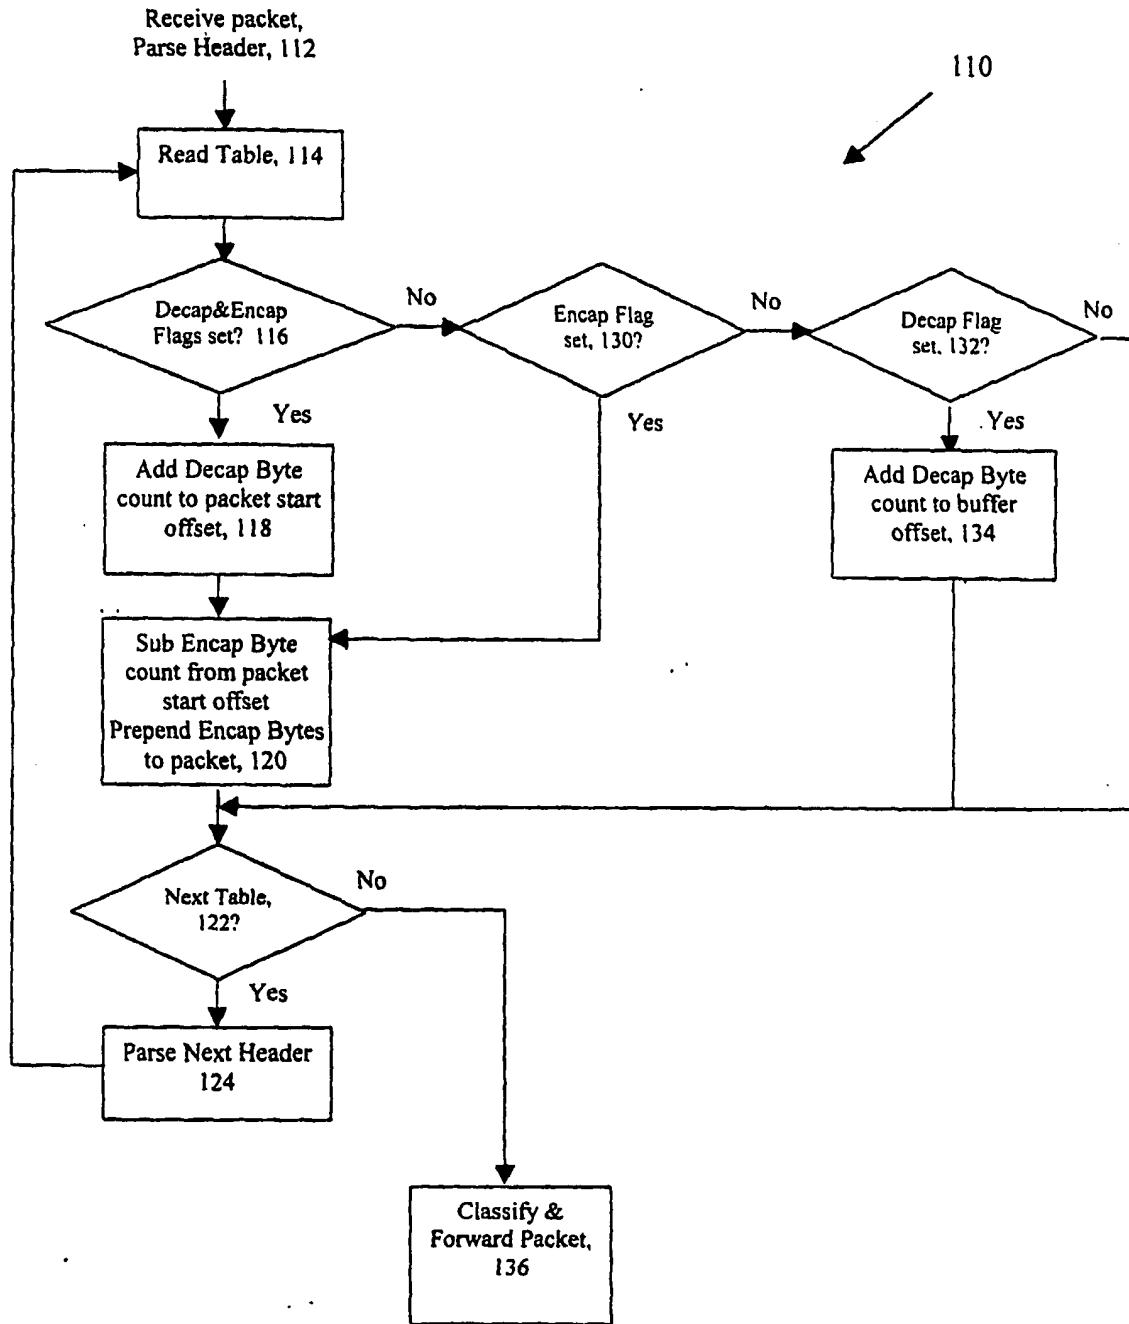


FIG. 6

9/9

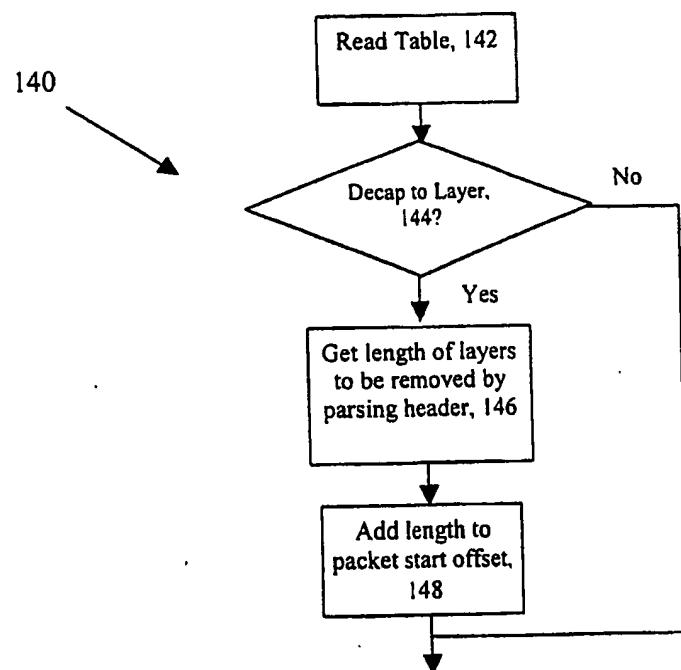


FIG. 7

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 02/29074

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 H04L12/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, PAJ, WPI Data, INSPEC, COMPENDEX, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>WO 98 20647 A (INTEGRATED TELECOM TECHNOLOGY) 14 May 1998 (1998-05-14)</p> <p>abstract</p> <p>figures 3,5,6</p> <p>page 2, line 36 -page 3, line 10</p> <p>page 5, line 3-26</p> <p>page 6, line 22 -page 8, line 21</p> <p>---</p> <p>-/-</p>	1,2,5,7, 9,10, 12-14, 19,20, 22, 24-26, 28-31
Y		3,4,27

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the International filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the International filing date but later than the priority date claimed

T later document published after the International filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

& document member of the same patent family

Date of the actual completion of the international search

18 December 2002

Date of mailing of the international search report

24/01/2003

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel: (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Gabriel, C

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 02/29074

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 651 002 A (READY DAVID C ET AL) 22 July 1997 (1997-07-22) abstract figures 1,3,4A,4B,7 column 1, line 6-13 column 2, line 3-48 column 2, line 65 -column 3, line 20 column 5, line 65 -column 6, line 44 column 17, line 66 -column 18, line 25 column 29, line 44 -column 30, line 23	1,19,28
Y	---	27
Y	EP 0 959 602 A (INST NAT RECH INF AUTOMAT) 24 November 1999 (1999-11-24) abstract figures 3A,3B,6,7 paragraphs '0063!-'0073!, '0078!, '0086! ---	3,4
X	WO 01 17179 A (IBM) 8 March 2001 (2001-03-08) abstract page 3, line 14 -page 4, line 9 page 6, line 19 -page 7, line 31 page 9, line 7-23 page 11, line 10 -page 13, line 10 page 14, line 1-15	1,2,10, 19,20, 22,28-30
A	---	3-9, 11-18, 21, 23-27,31
X	US 5 251 205 A (CALLON ROSS W ET AL) 5 October 1993 (1993-10-05) figures 5A,6 column 23, line 53-61 column 28, line 25-62 ---	1,2,19, 20,28,29

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 02/29074

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
WO 9820647	A 14-05-1998	WO	9820647 A1	14-05-1998
US 5651002	A 22-07-1997	NONE		
EP 0959602	A 24-11-1999	FR EP	2778809 A1 0959602 A1	19-11-1999 24-11-1999
WO 0117179	A 08-03-2001	US CN EP US US WO	6404752 B1 1369160 T 1208676 A1 2002061022 A1 2002048270 A1 0117179 A1	11-06-2002 11-09-2002 29-05-2002 23-05-2002 25-04-2002 08-03-2001
US 5251205	A 05-10-1993	US US	5430727 A 5557745 A	04-07-1995 17-09-1996